Reflectance Field Rendering of Human Faces for "Spider-Man 2"

Mark Sagar
Sony Pictures Imageworks

The creation of convincing computer generated human faces which can withstand close-up scrutiny under arbitrary lighting has been notoriously difficult to achieve, especially for well known actors. For the film "Spider-Man 2" it was decided to try recent experimental computer graphics research that looked promising but that had never been production tested.

In order to create highly realistic digital versions of the faces for the main characters (played by Tobey Maguire and Alfred Molina) the techniques introduced in the SIGGRAPH 2000 paper "Acquiring the reflectance field of a human face"[1] were extended to support deforming shapes for facial animation and integrated into the Sony Imageworks CG pipeline.

The USC ICT Lightstage 2.0 was used for the capture. Four film cameras were placed at various angles around the subject and synchronized to the strobes for simultaneous image capture. The resulting images were color corrected and projected onto a 3D model of the performer. Colorspace Analysis was used to decompose each dataset into diffuse and specular components. The specular component distribution was transformed to a view independent representation. The decomposed multiview datasets were then combined in UV space weighted by surface normal and visibility criteria to create the final reflectance functions.

In order to apply facial motion capture to create an animated model which could be relighted it is necessary either to capture images of the face in many expressions and interpolate, requiring a huge amount of data, or to modify a single set of neutral expression reflectance functions as the model deforms.
Because we had very limited time with the actors, and to avoid the logistics and processing cost of capturing many separate expressions, only the images captured for the faces in the neutral position (and with teeth bared) were used. In order to approximate the lighting changes as the face deformed (driven by motion capture or animation) into different expressions the reflectance functions were transformed based on changes in surface normal direction and light source visibility. This proved successful, however it is a first order approximation only as it does not account for indirect illumination changes and does not represent visual changes due to blood redistribution as the face compresses.

The eyes had to be treated separately as they move independently from the head, and the captured data for corneal highlights were too sparse to re-synthesize sharp continuous highlights. Obscured regions of the eyes were filled in by mapping normalized reflectance functions of nearby tissue. The surface geometry of the corneas was determined from the reflectance data, to create a bump map for a conventional shader.

A few unexpected problems had to be dealt with. The actors both had haircuts which did not match the final look in the film, and also had some surface blemishes. Certain small areas in the ear were not seen by any camera and so had no capture data. Several approaches used to fix these problems (which had to be addressed on 480 frames per camera) including paint fixes and the creation of a "virtual lightstage" to generate skin texture for missing regions.

One of the main challenges was integration with the conventional CG shading pipeline. To do this the reflectance field shader was implemented to work with both individual lights controlled by the existing lighting tools and with HDR image based environmental lighting. A downsampled dataset was available to speed up rendering for the environmental case. The acquired reflectance field is non-local so the data does not account for partial shadowing and other local lighting effects (e.g. the lighting changes caused by wearing sunglasses). This was approximated by locally modifying the illumination model using conventional shading techniques.

The successful combination of the reflectance field shaded and conventionally shaded CG permitted artists/Technical Directors to use preexisting tools and workflow without special consideration to which shading technique was being employed.

References
1. Debevec, P., Hawkins, T., Tchou, C., Duiker, H.P., Sarokin, W. and Sagar, M. 2000. Acquiring the reflectance field of a human face. In SIGGRAPH 2000, *Computer Graphics Proceedings*, 145-156, ACM SIGGRAPH

Figure 1 Digital vs Real